



Subscribe Register
(Full Service) (Limited Service, Free)

Login

ACM
Search

Search: The ACM Digital Library The Guide

"persistent url" purl "uniform resource name" urn

THE ACM DIGITAL LIBRARY

i F

Published before December 1998

Terms used persistent url purl uniform resource name urn

Sort results
by

relevance ▾

Save results to a Binder
 Search Tips
 Open results in a new window

Display results

expanded form ▾

Results 1 - 11 of 11

- 1 Location-independent naming for virtual distributed software repository**
Shirley Browne, Jack Dongarra, Stan Green, Keith Moore, Theresa Pepin, T
August 1995 **ACM SIGSOFT Software Engineering Notes**, Proceedings of the
Volume 20 Issue SI

Full text available: pdf(894.55 KB)

Additional Information: full citation, abstract, refere

A location-independent naming system for network resources has been d
description of software components accessible through a virtual distribut
easy and efficient searching and retrieval, and it addresses many of the
involved with distributed software repositories by providing mechanisms
and integrity checking. This paper ...

- 2 Resource discovery protocol for mobile computing**

Charles E. Perkins, Harry Harjono

December 1996

Mobile Networks and Applications, Volume 1 Issu

Full text available: pdf(246.43 KB)

Additional Information: full citation, abstract, refere

The increasing complexity of modern networks prompts a need for dyna
the additional need to rediscover the location of local area network resou
We present a protocol and proposal for the operation of dynamic resourc
and light weight. We implemented and tested our design with stationary
IP.

1

Location-independent naming for virtual distributed software repositories

Shirley Browne, Jack Dongarra, Stan Green, Keith Moore, Theresa Pepin, Tom Rowan, Reed Wade

August 1995

ACM SIGSOFT Software Engineering Notes , Proceedings of the 1995 Symposium on Software reusability,
Volume 20 Issue SI

Full text available: pdf(894.55 KB)

Additional Information: full citation, abstract, references, citings, index terms

A location-independent naming system for network resources has been designed to facilitate organization
and

description of software components accessible through a virtual distributed repository. This naming system
enables easy and efficient searching and retrieval, and it addresses many of the consistency, authenticity,
and

integrity issues involved with distributed software repositories by providing mechanisms for grouping
resources

and for authenticity and integrity checking. This paper ...

2

Resource discovery protocol for mobile computing

Charles E. Perkins, Harry Harjono

December 1996

Mobile Networks and Applications, Volume 1 Issue 4

Full text available: pdf(246.43 KB)

Additional Information: full citation, abstract, references, citings, index terms

The increasing complexity of modern networks prompts a need for dynamic resource discovery. Mobile
clients

have the additional need to rediscover the location of local area network resources each time they move to a
different LAN. We present a protocol and proposal for the operation of dynamic resource discovery. Our
design is

simple, extensible, and light weight. We implemented and tested our design with stationary servers, and
mobile

clients running mobile IP.

3

Scalable Networked Information Processing Environment (SNIPE)

Graham E Fagg, Keith Moore, Jack J Dongarra, Al Geist

November 1997

Proceedings of the 1997 ACM/IEEE conference on Supercomputing (CDROM)

Full text available: pdf(77.42 KB)

Additional Information: full citation, abstract, references, citings

SNIPE is a metacomputing system that aims to provide a reliable, secure, fault-tolerant environment for
long-term distributed computing applications and data stores across the global InterNet. This system
combines

global naming and replication of both processing and data to support large scale information processing
applications leading to better availability and reliability than currently available with typical cluster
computing

and/or distributed computer environments.

Keywords: MetaComputing, RCDS, SNIPE, reliable, scalable, secure

4

Position papers Supporting longevity in an information infrastructure architecture

Christos A. Papachriston

September 1996

X
Proceedings of the 7th workshop on ACM SIGOPS European workshop: Systems support for worldwide applications

Full text available: pdf(637.52 KB)

Additional Information: full citation, abstract, references

With the growth in the Internet and network-based community, comes an interest in building applications independently of the information or data on which they will operate. In conjunction with this divergence of information from application and significant growth in the amount of information available on the net the issues

of longevity are taking a more prominent position. As investment in information grows, the amount of information that has long-term value also grows. The URLs (Uniform Resourc ...

5

Types and persistence in database programming languages

Malcolm P. Atkinson, O. Peter Buneman

June 1987

ACM Computing Surveys (CSUR), Volume 19 Issue 2

Full text available: pdf(7.91 MB)

Additional Information: full citation, abstract, references, citings, index terms, review

Traditionally, the interface between a programming language and a database has either been through a set of relatively low-level subroutine calls, or it has required some form of embedding of one language in another. Recently, the necessity of integrating database and programming language techniques has received some long-overdue recognition. In response, a number of attempts have been made to construct programming languages with completely integrated database management systems. These lang ...

6

A case for document management functions on the Web

Gail L. Rein, Daniel L. McCue, Judith A. Stein

September 1997

Communications of the ACM, Volume 40 Issue 9

Full text available: pdf(415.20 KB)

Additional Information: full citation, references, citings, index terms

7

Referential integrity of links in open hypermedia systems

Hugh C. Davis

May 1998

X
Proceedings of the ninth ACM conference on Hypertext and hypermedia : links, objects, time and space---structure in hypermedia systems: links, objects, time and space---structure in hypermedia systems

Full text available: pdf(1.30 MB)

Additional Information: full citation, references, citings, index terms

8

Morphology & tagging: Morphology with a null-interface

Harald Trost, Johannes Matiasek

August 1994

Proceedings of the 15th conference on Computational linguistics - Volume 1

Full text available: [pdf\(553.76 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

We present an integrated architecture for word-level and sentence-level processing in a unification-based paradigm. The core of the system is a CLP implementation of a unification engine for feature structures supporting relational values. In this framework an HPSG-style grammar is implemented. Word-level processing

uses X2MORF, a morphological component based on an extended version of two-level morphology. This component is tightly integrated with the grammar as a relation. The advanta ...

9

Pearl: a probabilistic chart parser

David M. Magerman, Mitchell P. Marcus

April 1991

Proceedings of the fifth conference on European chapter of the Association for Computational Linguistics

Full text available: [pdf\(622.92 KB\)](#) [Publisher Site](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

This paper describes a natural language parsing algorithm for unrestricted text which uses a probability-based

scoring function to select the "best" parse of a sentence. The parser, Pearl, is a time-asynchronous bottom-up chart parser with Earley-type top-down prediction which pursues the highest-scoring theory in the chart, where

the score of a theory represents the extent to which the context of the sentence predicts that interpretation.

This parser differs from previous attempts at sto ...

10

Session V: Flexible parsing

Phil Hayes, George Mouradian

June 1980

Proceedings of the 18th conference on Association for Computational Linguistics

Full text available: [pdf\(699.04 KB\)](#) [Publisher Site](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

When people use natural language in natural settings, they often use it ungrammatically, missing out or repeating words, breaking-off and restarting, speaking in fragments, etc., Their human listeners are usually able

to cope with these deviations with little difficulty. If a computer system wishes to accept natural language input

from its users on a routine basis, it must display a similar indifference. In this paper, we outline a set of parsing

flexibilities that such a system should provide. ...

Reuse library interoperability and the World Wide Web

Shirley V. Browne, James W. Moore

May 1997

ACM SIGSOFT Software Engineering Notes , Proceedings of the 1997 symposium on Software reusability,

Volume 22 Issue 3

Full text available: pdf(1.16 MB)

Additional Information: full citation, references, citings, index terms


[Subscribe](#)
[Register](#)
[Login](#)

(Full Service) (Limited Service, Free)

Search: The ACM Digital Library The Guide

THE ACM DIGITAL LIBRARY
[Feedback Report](#)

Resource discovery protocol for mobile computing

[Full text](#)

Pdf (246 KB)

[Source](#)
Mobile Networks and Application
[Authors](#)

Volume 1 , Issue 4 (December 1996) tab

[Publisher](#)

Special issue on mobile computing and s

Pages: 447 - 455

Year of Publication: 1996

ISSN:1383-469X

Additional Information: abstract references citings index terms collaborative collea

Tools and Actions:
[Discussions](#)
[Find similar Articles](#)
[Review this Article](#)
[Save this Article to a Binder](#)
[Display in BibTex Format](#)

ABSTRACT

The increasing complexity of modern networks prompts a need for dynamic resou the additional need to rediscover the location of local area network resources each LAN. We present a protocol and proposal for the operation of dynamic resource di extensible, and light weight. We implemented and tested our design with stationai running mobile IP.

REFERENCES

Note: OCR errors may be found in this Reference List extracted from the full text article. ACM has o rather than only correct and linked references.

- 1 S. Alexander and R. Droms, DHCP options and BOOTP vendor extensions, RFC 1533 (October 1993) 9/22/04 12:15 PM

- 2 T. Berners-Lee, Universal resource identifiers in WWW, RFC 1630 (November 1994).
- 3 T. Berners-Lee, L. Masinter and M. McCahill, Uniform resource locators (URL), RFC1738 (December 1994).
- 4 C. Mic Bowman, P.B. Danzig, D.R. Hardy, U. Manber, M.F. Schwartz and D.P. Wessels, Harvest: access system, Technical Report CU- CS-732-94, University of Colorado, Computer Science Department (August 1994, Revised March 1995).
- 5 R. Caceres and L. Iftode, Improving the performance of reliable transport protocols in mobile computer Areas Commun. 13 (1995) 850-857.
- 6 Salutation Consortium, Salutation Architecture v2.0 , <http://www.salutation.org> (1996).
- 7 S. Deering and R. Hinden, Internet protocol, version 6 (IPv6) specification, RFC1883 (December 1995).
- 8 R. Droms, Dynamic host configuration protocol, RFC 1541 (October 1993).
- 9 P.E. Hoffman and R. Daniel, Jr., Generic URN syntax, draft-ietf-uri-urn-syntax-00.txt, work in progress.
- 10 M. Madsen, A critique of existing URN proposals, draft-ietfuri-urn-madsen-critique-00.txt, work in progress.
- 11 U. Manber and S. Wu, GLIMPSE: A tool to search through entire file systems, Technical Report Computer Science Department, Tucson, Arizona (October 1993).
- 12 C. E. Perkins , T. Jagannadh, DHCP for mobile networking with TCP/IP, Proceedings of the IEEE Communications (ISCC'95), p.255, June 27-29, 1995
- 13 C. Perkins (ed.), IPv4 mobility support, RFC 2002 (October 1996).
- 14 B. Pinkerton, Finding what people want: Experiences with the WebCrawler, Proc. Second International Conference on the Web , Chicago (October 1994).
- 15 B. Schilit and M. Theimer, Disseminating active map information to mobile hosts, IEEE Network 9(5) (1995) 12-20.
- 16 K. Shafer, E. Miller, V. Tkac and S. Weibel, URN services, draft-shafer-uri-urn-resolution-00.txt, work in progress.
- 17 K. Sollins and L. Masinter, Functional requirements for uniformresource names, RFC 1737 (December 1994).
- 18 J. Veizades, E. Guttman, C. Perkins and S. Kaplan, Service location protocol, draft-ietf-srvloc-p01.txt, work in progress.
- 19 Roy Want , Andy Hopper , Veronica Falcão , Jonathan Gibbons, The active badge location system, TOIS Systems (TOIS), v.10 n.1, p.91-102, Jan. 1992
- 20 S. Wu and U. Manber, A fast algorithm for multi-pattern searching, Technical Report TR 94-17, University of Southern California, Los Angeles, CA, USA.
- 21 A. Young, Connection-less lightweight directory access protocol, RFC 1798 (June 1995).

↑ CITINGS

Mike Esler , Jeffrey Hightower , Tom Anderson , Gaetano Borriello, Next century computing: the Portolano project at the University of Washington, Proc. annual ACM/IEEE international conference on Mobile computing and networking, pp. 1-10, Seattle, Washington, United States

↑ INDEX TERMS

Primary Classification:

C. Computer Systems Organization

↳ C.2 COMPUTER-COMMUNICATION NETWORKS

↳ C.2.1 Network Architecture and Design

↳ Subjects: Wireless communication

Additional Classification:

C. Computer Systems Organization

↳ C.2 COMPUTER-COMMUNICATION NETWORKS

↳ C.2.2 Network Protocols

↳ Subjects: Applications (SMTP, FTP, etc.)

↳ Nouns: IP

H. Information Systems

↳ H.3 INFORMATION STORAGE AND RETRIEVAL

↳ H.3.3 Information Search and Retrieval

↳ Subjects: Search process

General Terms:

Algorithms, Design

↑ Collaborative Colleagues:

Harry Harjono: Charles E. Perkins

Charles E. Perkins:	Sherman R. Alpert Elizabeth M. Belding-Royer Pravin Bhagwat Imrich Chlamtac Petre Dini J. J. Garcia-Luna-Aceves Harry Harjono David B. Johnson Birgitta König-Ries Lev Kacnelson	Sung-Ju Lee Kevin Luo Kia Makki S. A. M. Makki Jari T. Malinen Marc Mosko Anders Nilsson Niki Pissinou Peter L. Reiher Elizabeth M. Royer	Peter Scheuermann Burkhard Stiller C-K. Toh Antti J. Tuominen Nitin H. Vaidya Jari Veijalainen Ryuji Wakikawa Alexander L. Wolf Ouri Wolfson Bobby Woolf
---------------------	---	--	---

↑ Peer to Peer - Readers of this Article have also read:

- Web application security assessment by fault injection and behavior modeling
Proceedings of the twelfth international conference on World Wide Web
Yao-Wen Huang , Shih-Kun Huang , Tsung-Po Lin , Chung-Hung Tsai
- Data structures for quadtree approximation and compression
Communications of the ACM 28, 9
Hanan Samet
- A hierarchical single-key-lock access control using the Chinese remainder theorem
Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computation
Kim S. Lee , Huizhu Lu , D. D. Fisher
- 3D representations for software visualization
Proceedings of the 2003 ACM symposium on Software visualization
Andrian Marcus , Louis Feng , Jonathan I. Maletic

- Probabilistic surfaces: point based primitives to show surface uncertainty
Proceedings of the conference on Visualization '02
Gevorg Grigoryan , Penny Rheingans

The ACM Portal is published by the Association for Computing Machinery. Copyright
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)


 Search: The ACM Digital Library The Guide

THE ACM DIGITAL LIBRARY
[Feedback Report](#)

Supporting personal mobility for nomadic computing over the internet

Full text
[Pdf \(1.42 MB\)](#)
Source
**ACM SIGMOBILE Mobile Com
Review archive**

Volume 1 , Issue 1 (April 1997) table

Pages: 22 - 31

Year of Publication: 1997

Authors

 Yalun Li The University
Canada

 Victor C. M. Leung The University
Canada

Publisher

ACM Press New York, NY, US

Additional Information: [abstract](#) [references](#) [citations](#) [collaborative colleagues](#)
Tools and Actions:
[Discussions](#)
[Find similar Articles](#)
[Review this Article](#)
[Save this Article to a Binder](#)
[Display in BibTex Format](#)
DOI Bookmark:

 Use this link to bookmark this Article: <http://doi.acm.org/10.1145/583973.5>
[What is a DOI?](#)

↑ ABSTRACT

This paper presents a new paradigm for nomadic computing over the Internet called (UPC), where mobile users can access computing resources, network services, and environments anywhere using any available terminals. The concept of UPC and system discussed, and the required system architecture capable of managing different mobile terminals, in the UPC environment is presented. Modifications of connection setup application programs to enable addressing based on a global user identity are considered.

Note: OCR errors may be found in this Reference List extracted from the full text ;
the complete List rather than only correct and linked references.

- 1 Leonard Kleinrock, Nomadic computing—an opportunity, ACM SIGCOMM Review, v.25 n.1, p.36-40, Jan. 1995
- 2 R. Bagrodia, W.W. Chu, L. Kleinrock, and G. Popek, "Vision, Issues, and Architectures for Personal Communications," IEEE Personal Communications Magazine, pp. 14--27, Dec. 1995.
- 3 W. Stallings, "IPv6: The New Internet Protocol," IEEE Communications Magazine, vol. 33, pp. 12-17, June 1995.
- 4 C. Perkins, editor. RFC 2002, IP Mobility Support, October 1996.
- 5 CCITT SG 1, Draft Recommendation F.851 UPT-service Description, Version 7.
- 6 G. S. Lauer, "IN Architectures for Implementing Universal Personal Telecommunications," IEEE Communications Magazine, vol. 8, pp. 6--16, Mar./Apr. 1994.
- 7 M. Fujioka, S. Sakai, H. Yagi, "Hierarchical and Distributed Information Handling," IEEE Communications Magazine, vol. 4, pp. 50--60, Nov. 1990.
- 8 J. Vandenameele, J. B. Thieffry, H. Decuyper, "UPT, A New Dimension in Telecommunications," Proc. Digital Communications Zurich Seminar, pp. 41--54, 1992.
- 9 ETSI DTR/NA-70201, UPT: General Service Description, Sept. 1991.
- 10 F. Dupuy, G. Nilsson, and Y. Inoue, "The TINA Consortium: Toward Networked Information Services," IEEE Communications Magazine, pp. 78--83, Nov. 1995.
- 11 T. Eckardt and T. Magedanz, "From IN towards TINA-based Personal Communications," Proc. ICUPC '92, pp. 402--406, Sept. 1996.
- 12 P. Bhagwat, C. Perkins, and S. Tripathi, "Network Layer Mobility: An Architecture," IEEE Communications Magazine, pp. 54--64, June 1996.
- 13 A. Myles, D. B. Johnson, and C. Perkins, "A Mobile Host Protocol Supporting Remote Authentication," IEEE Journal on Selected Areas in Communications, vol. 13, pp. 821-834, Oct. 1995.
- 14 A. Nakajima, M. Kuramoto, and M. Yabusaki, "Intelligent Network Architecture for Telecommunication," Proc. ICUPC '92, pp.339--344, 1992.
- 15 S. Mohan and R. Jain, "Two User Location Strategies for Personal Communications," IEEE Personal Communications Magazine, pp. 42--50, First Quarter 1994.
- 16 Douglas E. Comer, Internetworking with TCP/IP (2nd ed.), vol. I, Prentice-Hall, Inc., 1991
- 17 D. Chess, et. al., "Itinerant Agents for Mobile Computing," IEEE Personal Communications Magazine, pp. 34--49, Oct. 1995.
- 18 I. Iida, T. Nishigaya, and K. Murakami, "DUET: An Agent-Based Personal Communication System," Proc. ICUPC '92, pp. 353--357, 1992.

- Communications Magazine, pp. 44-49, Nov. 1995.
- 19 R. Ramjee, T. F. L. Porta, and M Veeraraghavan, "The Use of Network-Based Personal Communication Services," IEEE Personal Communications Magazine, pp.
- 20 K. M. Chandy , A. Rifkin , P. A. G. Sivilotti , J. Mandelson , M. Richardson , W. world-wide distributed system using Java and the Internet, Proceedings of the Hig Computing (HPDC '96), p.11, August 06-09, 1996

↑ CITINGS

Mária Törö , Thong Tri Huynh , Jinsong Zhu , Kangming Liu , Victor C. M. Leung, CORBA based des personal computing, Mobile Networks and Applications, v.8 n.1, p.75-86, February 2003

↑ Collaborative Colleagues:

Victor C. M. Leung:	Andrew W. Y. Au	Mária Törö
	Henry C. B. Chan	Maria Törö
	Jean-François Frigon	Son T. Vuong
	Sergio Gonzalez-Valenzuela	Vincent W. S. Wong
	Thong Tri Huynh	W. S. Vincent Wong
	Chris Jedrzycki	Oliver T. W. Yu
	Mark E. Lewis	Jinsong Zhu
	Yalun Li	
	Kangming Liu	
	Fei Peng	
Yalun Li:	Victor C. M. Leung	

The ACM Portal is published by the Association for Computing Machinery. Copyright ©
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)


 Search: The ACM Digital Library The Guide

mobile <and> computing

THE ACM DIGITAL LIBRARY

Fe

Terms used mobile and computing

 Sort results
by

 relevance

 Save results to a Binder

Display results

 expanded form

 Search Tips

 Open results in a new window

Results 141 - 160 of 200

Result page: previous 1 2 3 4 5 6

Best 200 shown

141 Intrusion detection techniques for mobile wireless networks

Yongguang Zhang, Wenke Lee, Yi-An Huang

September 2003

Wireless Networks, Volume 9 Issue 5

Full text available: pdf(164.73 KB)

Additional Information: full citation, abstract, refe

The rapid proliferation of wireless networks and mobile computing application security. The traditional way of protecting networks with firewalls and encryption is effective. We need to search for new architecture and mechanisms to protect mobile computing application. In this paper, we examine the vulnerabilities of wireless intrusion detection in the security ...

Keywords: anomaly detection, cooperative detection, intrusion detection, intre

142 Dynamic task-based anycasting in mobile ad hoc networks

Prithwish Basu, Wang Ke, Thomas D. C. Little

October 2003

Mobile Networks and Applications, Volume 8 Issue 5

Full text available: pdf(518.68 KB)

Additional Information: full citation, abstract, refe

Mobile ad hoc networks (MANETs) have received significant attention in the recent years due to the increasing numbers of tetherless portable devices, and rapid growth in popularity of wireless communication technologies. The research community has remained focused on developing lower layer mechanisms and protocols for making MANETs operational. However, little focus has been applied on high level issues such as distributed application modeling and management in dynamic MANET environments. ...

Keywords: anycasting, device/service discovery, distributed application execu

143 Foundation of a framework to support knowledge management in the field of computing

Philipp Amann, Gerald Quirchmayr

January 2003 Proceedings of the Australasian information security workshop conf

Full text available: [pdf\(761.39 KB\)](#)

Additional Information: full citation, abstract, references

In this paper we propose a framework to combine Knowledge Management and computing emphasizing on synchronization and adaptation issues of workflow processes. The proposed framework is to enable adaptive, two-way interaction between context and settings. In contrast to existing concepts, we aim at capturing active feedback from the environment. *Organizational Memory*, after ...

Keywords: WfMS, adaptability, context-awareness, knowledge management, synchronization

144 Power conservation strategy for mobile computers using load sharing

Mazliza Othman, Stephen Hailes

January 1998 ACM SIGMOBILE Mobile Computing and Communications Review

Full text available: [pdf\(1.10 MB\)](#)

Additional Information: full citation, abstract, references

Power management is an important aspect of mobile computing. Previous work has concentrated on the hardware approach. In this paper, we propose a different approach for mobile computers which is based on the concept of load sharing. User job can be moved from fixed host to reduce power consumption by the CPU. Simulation results show that this approach can extend battery lifetime by up to 2 ...

145 Using a coordination language to specify and analyze systems containing mobile entities

P. Ciancarini, F. Franzé, C. Mascolo

April 2000 ACM Transactions on Software Engineering and Methodology (TOSEM)

Full text available: [pdf\(306.92 KB\)](#)

Additional Information: full citation, abstract, references

New computing paradigms for network-aware applications need specification languages for mobile code-based systems. A coordination language provides a formal framework for specifying the behavior of mobile entities. Entities can be expressed. A coordination language deals with the creation and destruction of entities, their communication activities, as well as their distribution and mobility in space. The PoliS language offers a flexible ...

146 Composable ad hoc location-based services for heterogeneous mobile environments

Todd D. Hodes, Randy H. Katz

October 1999 Wireless Networks, Volume 5 Issue 5

Full text available: [pdf\(403.18 KB\)](#)

Additional Information: full citation, references, citations, index terms

147 Modeling mobile IP in mobile UNITY

Peter J. McCann, Gruia-Catalin Roman

April 1999 ACM Transactions on Software Engineering and Methodology (TOSEM)

Full text available:  pdf(344.70 KB)

Additional Information: full citation, abstract, references,

With recent advances in wireless communication technology, mobile computing research. A mobile system is one where independently executing component: the course of the computation, and where the pattern of connectivity among and out of proximity. Mobile UNITY is a notation and proof logic for specifying this article it is argued that Mobile ...

Keywords: formal methods, mobile UNITY, mobile computing, shared variable, weak consistency

148 Cyberguide: a mobile context-aware tour guide

Gregory D. Abowd, Christopher G. Atkeson, Jason Hong, Sue Long, Rob Kooper
October 1997 Wireless Networks, Volume 3 Issue 5

Full text available:  pdf(596.81 KB)

Additional Information: full citation, abstract, references

Future computing environments will free the user from the constraints of the environment should take advantage of contextual information, such as position. In this paper, we present the Cyberguide project, in which we are building prototypes. Knowledge of the user's current location, as well as a history of past locations: services that we co ...

149 Models for mobile computing agents

Sanjiva Prasad

December 1996 ACM Computing Surveys (CSUR)

Full text available:  html(3.92 KB) Additional Information: full citation, citations

150 Design and analysis of an efficient communication strategy for hierarchical networks

Ioannis Chatzigiannakis, Sotiris Nikoletseas

August 2004

Mobile Networks and Applications, Volume 9 Issue 4

Full text available: [pdf\(313.32 KB\)](#)

Additional Information: full citation, abstract, refer

In this work we introduce two practical and interesting models for hierarchical ad-hoc networks, comprised of dense subnetworks supported by a very fast yet limited backbone infrastructure, (b) highly dynamic networks. If the deployment area changes in a highly dynamic way and is composed of several small clusters, we study the problem of basic communication, i.e., from a source node to a receiver node. ...

Keywords: ad-hoc mobile networks, average case analysis, expected walks, routing

151 New directions: A message ferrying approach for data delivery in sparse MANETs

W. Zhao, M. Ammar, E. Zegura

May 2004 Proceedings of the 5th ACM international symposium on Mobile ad hoc networking and computing

Full text available: [pdf\(152.16 KB\)](#)

Additional Information: full citation, abstract, refer

Mobile Ad Hoc Networks (MANETs) provide rapidly deployable and self-configurable solutions for critical applications, e.g., battlefields, disaster relief and wide area sensing. In this paper, we propose a new direction for efficient data delivery in *sparse* MANETs where network partitions can last for hours. We propose a message ferrying scheme that does not rely on the use of either long range communication which leads to rapid drain of battery power or frequent node mobility which leads to high loss rates. ...

Keywords: mobile ad hoc networks, mobility-assisted data delivery, sparse networks

152 Challenges: an application model for pervasive computing

Guruduth Banavar, James Beck, Eugene Gluzberg, Jonathan Munson, Jeremy Sussman

August 2000 Proceedings of the 6th annual international conference on Mobile computing and networking

Full text available: [pdf\(998.42 KB\)](#)

Additional Information: full citation, abstract, references

The way mobile computing devices and applications are developed, deployed and used fails to meet the expectations of the user community and falls far short of the potential for personal computing. This paper challenges the mobile computing community by questioning the roles of devices, applications and environments in pervasive computing. A new application model for pervasive computing is described, along with attributes of a new application model and challenges that must be met. ...

Olaf Zukunft

April 1997 Proceedings of the 1997 ACM symposium on Applied computing

Full text available:  pdf(770.76 KB) Additional Information: full citation, references, citings, index terms

Keywords: active database, configuration, mobile computing, mobile databases

154 Integrated network computing models, programming modes and software

V. K. Murthy, E. V. Krishnamurthy

February 1998 Proceedings of the 1998 ACM symposium on Applied Computing

Full text available:  pdf(698.60 KB) Additional Information: full citation, references, index terms

155 Composable ad-hoc mobile services for universal interaction

Todd D. Hodes, Randy H. Katz, Edouard Servan-Schreiber, Lawrence Rowe

September 1997 Proceedings of the 3rd annual ACM/IEEE international conference

Full text available:  pdf(1.86 MB) Additional Information: full citation, references, citing

156 Untraceability in mobile networks

Didier Samfat, Refik Molva, N. Asokan

December 1995 Proceedings of the 1st annual international conference on Mobile

Full text available:  pdf(1.20 MB) Additional Information: full citation, references, citings, index terms

Keywords: CDPD, GSM, alias, anonymity, authentication, location privacy, mobile networks

157 Using publish/subscribe middleware for mobile systems

Gianpaolo Cugola, H.-Arno Jacobsen

October 2002 ACM SIGMOBILE Mobile Computing and Communications Review

Full text available:  pdf(92.71 KB) Additional Information: full citation, abstract, reference

The range of mobile computing applications comprises location-based service. Middleware for these applications must effectively support the interaction of ; decoupled processing, and mediate between potentially millions of mobile clients achieve with traditional client/server middleware systems. We argue that the system addresses many of the challenge ...

168 Application requirements for middleware for mobile and pervasive system

Kimmo Raatikainen, Henrik Bærbak Christensen, Tatsuo Nakajima

October 2002 ACM SIGMOBILE Mobile Computing and Communications Review

Full text available: pdf(111.22 KB)

Additional Information: full citation, abstract, references

In this paper, we examine the requirements for future middleware to support identify key research areas. We illustrate the research areas with requirement projects concerning pervasive healthcare and home entertainment.

159 Multi-fidelity algorithms for interactive mobile applications

M. Satyanarayanan, Dushyanth Narayanan

November 2001

Wireless Networks, Volume 7 Issue 6

Full text available: pdf(111.52 KB)

Additional Information: full citation, abstract, references

We introduce the concept of *multi-fidelity algorithms*, which revises the class having a fixed output criterion and allowing the resource consumption to vary, to allow the *fidelity* or output criterion to vary. We discuss how multi-fidelity algorithms can extend battery life of interactive mobile applications. An extension of this idea allows *spots*: ...

Keywords: *Aura, Odyssey, adaptation, augmented reality, mobile computing, spot*

160 A framework for delivering multicast message in networks with mobile hosts

Arup Acharya, B. R. Badrinath

October 1996

Mobile Networks and Applications, Volume 1 Issue 2

Full text available: pdf(438.33 KB)

Additional Information: full citation, abstract, references

To accommodate mobile hosts (MHs) within existing data networks, the static support stations (MSSs) that communicate directly with MHs, usually make overall network changes dynamically as MHs connect to the static network from different locations at different times. Compared to their desktop counterparts, mobile hosts face the bandwidth of the wireless links ...

Results 141 - 160 of 200

Result page: previous 1 2 3

The ACM Portal is published by the Association for Computing Machinery

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: Adobe Acrobat QuickTime Windows Media

- 2 T. Berners-Lee, Universal resource identifiers in WWW, RFC 1630 (November 1994).
- 3 T. Berners-Lee, L. Masinter and M. McCahill, Uniform resource locators (URL), RFC1738 (December 1994).
- 4 C. Mic Bowman, P.B. Danzig, D.R. Hardy, U. Manber, M.F. Schwartz and D.P. Wessels, Harvest: access system, Technical Report CU- CS-732-94, University of Colorado, Computer Science Department (August 1994, Revised March 1995).
- 5 R. Caceres and L. Iftode, Improving the performance of reliable transport protocols in mobile computing environments, Areas Commun. 13 (1995) 850-857.
- 6 Salutation Consortium, Salutation Architecture v2.0 , <http://www.salutation.org> (1996).
- 7 S. Deering and R. Hinden, Internet protocol, version 6 (IPv6) specification, RFC1883 (December 1995).
- 8 R. Droms, Dynamic host configuration protocol, RFC 1541 (October 1993).
- 9 P.E. Hoffman and R. Daniel, Jr., Generic URN syntax, draft-ietf-uri-urn-syntax-00.txt, work in progress.
- 10 M. Madsen, A critique of existing URN proposals, draft-ietfuri-urn-madsen-critique-00.txt, work in progress.
- 11 U. Manber and S. Wu, GLIMPSE: A tool to search through entire file systems, Technical Report Computer Science Department, Tucson, Arizona (October 1993).
- 12 C. E. Perkins , T. Jagannadh, DHCP for mobile networking with TCP/IP, Proceedings of the IEEE Communications (ISCC'95), p.255, June 27-29, 1995
- 13 C. Perkins (ed.), IPv4 mobility support, RFC 2002 (October 1996).
- 14 B. Pinkerton, Finding what people want: Experiences with the WebCrawler, Proc. Second International Conference on the Web , Chicago (October 1994).
- 15 B. Schilit and M. Theimer, Disseminating active map information to mobile hosts, IEEE Network 9(5) (1995) 12-20.
- 16 K. Shafer, E. Miller, V. Tkac and S. Weibel, URN services, draft-shafer-uri-urn-resolution-00.txt, work in progress.
- 17 K. Sollins and L. Masinter, Functional requirements for uniformresource names, RFC 1737 (December 1994).
- 18 J. Veizades, E. Guttman, C. Perkins and S. Kaplan, Service location protocol, draft-ietf-svrlc-p00.txt, work in progress.
- 19 Roy Want , Andy Hopper , Veronica Falcão , Jonathan Gibbons, The active badge location system, TOIS, v.10 n.1, p.91-102, Jan. 1992
- 20 S. Wu and U. Manber, A fast algorithm for multi-pattern searching, Technical Report TR 94-17, University of Southern California, Los Angeles, CA, USA.
- 21 A. Young, Connection-less lightweight directory access protocol, RFC 1798 (June 1995).

↑ CITINGS

Mike Esler , Jeffrey Hightower , Tom Anderson , Gaetano Borriello, Next century computing: the Portolano project at the University of Washington, Annual ACM/IEEE international conference on Mobile computing and networking, pp. 1-10, Seattle, Washington, United States

↑ INDEX TERMS

Primary Classification:

C. Computer Systems Organization

↳ C.2 COMPUTER-COMMUNICATION NETWORKS

C.2.1 Network Architecture and Design

↳ Subjects: Wireless communication

Additional Classification:

C. Computer Systems Organization

↳ C.2 COMPUTER-COMMUNICATION NETWORKS

↳ C.2.2 Network Protocols

↳ Subjects: Applications (SMTP, FTP, etc.)

↳ Nouns: IP

H. Information Systems

↳ H.3 INFORMATION STORAGE AND RETRIEVAL

↳ H.3.3 Information Search and Retrieval

↳ Subjects: Search process

General Terms:

Algorithms, Design

↑ Collaborative Colleagues:

Harry Harjono: Charles E. Perkins

Charles E. Perkins:	Sherman R. Alpert Elizabeth M. Belding-Royer Pravin Bhagwat Imrich Chlamtac Petre Dini J. J. Garcia-Luna-Aceves Harry Harjono David B. Johnson Birgitta König-Ries Lev Kacnelson	Sung-Ju Lee Kevin Luo Kia Makki S. A. M. Makki Jari T. Malinen Marc Mosko Anders Nilsson Niki Pissinou Peter L. Reiher Elizabeth M. Royer	Peter Scheuermann Burkhard Stiller C-K. Toh Antti J. Tuominen Nitin H. Vaidya Jari Veijalainen Ryuji Wakikawa Alexander L. Wolf Ouri Wolfson Bobby Woolf
---------------------	---	--	---

↑ Peer to Peer - Readers of this Article have also read:

- Web application security assessment by fault injection and behavior modeling
Proceedings of the twelfth international conference on World Wide Web
Yao-Wen Huang , Shih-Kun Huang , Tsung-Po Lin , Chung-Hung Tsai
- Data structures for quadtree approximation and compression
Communications of the ACM 28, 9
Hanan Samet
- A hierarchical single-key-lock access control using the Chinese remainder theorem
Proceedings of the 1992 ACM/SIGAPP Symposium on Applied computation
Kim S. Lee , Huizhu Lu , D. D. Fisher
- 3D representations for software visualization
Proceedings of the 2003 ACM symposium on Software visualization
Andrian Marcus , Louis Feng , Jonathan I. Maletic

- Probabilistic surfaces: point based primitives to show surface uncertainty
Proceedings of the conference on Visualization '02
Gevorg Grigoryan , Penny Rheingans

The ACM Portal is published by the Association for Computing Machinery. Copyright
[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)

[Homepage](#) | [Advanced Search](#)Search using: [HotBot](#) [Google](#) [Ask Jeeves](#)

CUSTOM WEB FILTERS

[Tools](#) | [HotBot Skins](#) | [Preferences](#)Date: Before December 1 1998 [[Edit this Search](#)]WEB RESULTS by  (Showing Results 1 - 1 of 1)

1. CC: elmo@onelist.com Subject: Dataset servers Content-Type: text
...tcl/url.html and http://purl.org/thecliff/tcl/url This uses teh OCLC's "Persistent URL ... Each entry also has a "modification timestamp" as...
www.oclc.de/~akupries/step/bib/jcw.dataset.srv.txt

Search for ""persistent url" ~timestamp" using: [HotBot](#), [Google](#)

ask jeeves search (via hotbot)

[Advertise](#) | [Help](#) | [Text-only Skin](#) | [Submit Site](#) | [HotBot International](#) | [Yellow Pages](#)© Copyright 2004, Lycos, Inc. All Rights Reserved. | [Privacy Policy](#) | [Terms & Conditions](#) | [HotBot Your Site](#)

NOTE: You can't see most of our site's design details because your browser doesn't support basic [Web standards](#). You should consider upgrading to a [more recent browser release](#). If you choose to continue with the use of your current browser however, all of our content will continue to be accessible to all versions of every browser.

Return-Path: bluepeak.westend.com!popeye
Return-Path: <popeye@bluepeak.westend.com>
Received: from popeye.bluepeak.westend.com by bluepeak with smtp
(Smail3.2 #1) id m0zYbGa-00027DC; Wed, 28 Oct 1998 20:28:24 +0100 (MET)
Received: from genesis for a.kupries
with Cubic Circle's cucipop (v1.10 1996/09/06) Wed Oct 28 20:21:08 1998
X-From_: jcw@equi4.com Wed Oct 28 19:00:00 1998
Received: from vservers.com (root@[207.159.153.130] (may be forged))
by genesis.westend.com (8.8.6/8.8.6) with ESMTP id SAA06539
for <a.kupries@westend.com>; Wed, 28 Oct 1998 18:59:57 +0100 (MET)
Received: from mini.net ([207.159.134.5])
by vservers.com (8.9.0/8.9.0) with ESMTP id JAA13141;
Wed, 28 Oct 1998 09:57:52 -0800 (PST)
Received: from equi4.com (siepie.equi4.nl [195.108.246.51])
by mini.net (8.8.5/8.8.5) with ESMTP id PAA08975;
Wed, 28 Oct 1998 15:09:13 GMT
Message-ID: <363732E1.9C3BDBCE@equi4.com>
Date: Wed, 28 Oct 1998 16:06:09 +0100
From: Jean-Claude Wippler <jcw@equi4.com>
Organization: Equi4 Software - <http://www.equi4.com>
X-Mailer: Mozilla 4.5b2 [en] (Win98; I).
X-Accept-Language: en
MIME-Version: 1.0
To: Alexandre Ferrieux <alexandre.ferrieux@cnet.francetelecom.fr>,
Andreas Kupries <a.kupries@westend.com>,
Brent Welch <welch@scriptics.com>,
Cameron Laird <claird@Starbase.NeoSoft.COM>,
Larry Virden <lvirden@cas.org>, Mark Roseman <roseman@teamwave.com>
CC: elmo@onelist.com
Subject: Dataset servers
Content-Type: text/plain; charset=us-ascii
Content-Transfer-Encoding: 7bit

All,

This is a restarted attempt to describe some details of what seem to be falling into place, and what I am starting to call "dataset servers".

The first start I made yesterday ended in so much handwaving (so to speak), and so much head-in-the-clouds talk that I abandoned it...

Another attempt, with about one page of context and justification, then the details of what has been set up so far and a few immediate plans.

THE PROBLEM

To find a way to share lists of information, such as bookmarks, FAQs, and the items we've been focusing on: patches, bug lists, several other valuable lists incorporated in "Jumbo", and Tcl-URL!.

Each of these is a dataset, with an unambiguous name, and a way to identify entries in them. Nothing new here.

The problems with today's datasets are twofold:

- How to get more people involved to maintain them. Tcl-URL! is a recent example: though others are doing the big job of compiling content, there's overhead in maintaining the infrastructure, and as long as that's manual, it will prevent everyone from being more and more productive. Automation is the only way to scale up... Content is fun, but the luggage associated with it is not.
- How to get accurate, painstakingly maintained lists, such as FAQs to as many people as possible. Grabbing a copy solves a problem now, but creates a maintenance task for the future. And I *want*

a copy of lists such as big reports and FAQs - on my disk. Because there are very powerful ways to browse them (read: "incrFilter").

There's another aspect, which is closely related to both problems. The web is full of lists, in HTML. Generated from some database perhaps, even dynamically (example: Scriptics), but the data is no longer easily accessible as a table of records, with a key and modification times. That's fine for a few lines of text, just grab a new copy every once in a while, but it's very inconvenient for larger amounts of information.

When someone makes some changes to in the manual pages, I'd like to get those. Not once, not every 6 months, but *frequently*, and with as little effort as possible. For bugs and patches, I want a *very* up to date set to browse in. After all, some packages are actively updated. The "RTFM" phrase is a response to a symptom. The underlying problem is that documentation is a *big* hassle to track and maintain. Out of date documentation is worth so much less, that we all end up not using it.

TODAY'S SOLUTIONS

Go to the web. Always. Use the master copy and never anything else. Sure, spend a lot of time (waiting), and a lot of money (connecting)... Sorry, but this is *not* a solution for most lists. Not for me.

Check on Usenet, c.l.t for example. When in doubt, post a message. Again, that's not workable: Usenet is a volatile medium. Everything of value fades away within a week or so. The result, a range of frequently asked questions (and regulars doing an unbelievably good *manual* job of replying and pointing people to places where answers can be found).

Check DejaNews, the marriage of Usenet and the Web. Ahem, that's not too effective: how do you find an answer, even with a perfect search engine, in a database which is growing at the rate of a few *million* messages a week... (if I may believe the sequence numbers of DejaNews).

Spider a reference web site, and store a local copy. Yes, that's an option. It can take some time to do this, and to re-scan for changes. It also is likely to break down on dynamically generated pages. I've tried this a few times - but there is no site which covers exactly my interests and not too much more. And refresh really is slow...

Rely on peers, send each other tips and hints through email. Subscribe to mailing lists and archive all emails. The trouble with this, is that email archives are not really *that* useful when it's a collection of a few thousand emails, and nothing but plain text search as access method. There is too little focus, there is no organization. And it's a drag to carefully place each email in its proper spot (one?). Automated or not.

DATASET SERVERS

So, let's try something else. Let's make it easy for people to stay in control of the datasets they maintain (and own), while at the same time adding an infrastructure for others to get a copy, and to track changes. On the web, on servers, as databases, as HTML pages, but also as local data, available for use while off-line (most of the internet world is still very much affected by the distinction of on-line vs. off-line).

Let's give such datasets a clear identity, and let's create a mechanism whereby anyone can find the master copy, even if it moves occasionally. Let's do it in such a way that there is no central management / control.

Let's aim high, and make this approach future-proof. Usable in all sorts of contexts (including those that have not yet been created).

Let's keep our feet on the ground and start simple, with some very clear

and immediate goals. And let's keep focus on the benefits right away.

Let's treat organizations and individuals on the same level. Both as content providers, and as content consumers. Let the content be public, and available for use on websites, as well as on people's local disks.

Let's create "dataset servers".

AND THIS IS HOW

A dataset is a collection of entries (a table of records). Each dataset has a unique name. I'll use the Tcl-URL! posting archive as example, because it is furthest along in implementation.

The dataset name is: "tcl/url". It has two locations on the net associated to it:

<http://purl.org/thecliff/tcl/url.html>

and

<http://purl.org/thecliff/tcl/url>

This uses teh OCLC's "Persistent URL" service. OCLC is an independent organization which sees it as its task to support and enhance archival services. They have a commitment to be around and stay around. Their PURL service is a redirection service which re-routes HTTP requests. It's similar, but much simpler than the holy grail of archival: URNs. More importantly, it works. Today. And it scales well.

The URL "<http://purl.org/thecliff/tcl/url.html>" can be entered in a web browser, and leads to a page related to the Tcl-URL! archive dataset. Its contents and meaning is merely: "this is the home of Tcl-URL!". It can be anywhere - PURLs can be adjusted if it needs to be relocated any time in the future. PURLs can never go away. The worst that can happen, is that they are set to point to a "null URL". Then the page shown will be the history of where the PURL *used* to point to. That history is, just like the PURL itself, permanent. Available forever.

The URL "<http://purl.org/thecliff/tcl/url>" is the core of this system. It points to a "dataset server" which must adhere to a fixed protocol. Both PURLs can point to any web server, which need not even be the same.

There are two management issues for PURLs. They can be owned by one or more people - allowing them to alter the redirection, and there are PURL domain writers, which lets people create new PURLs (again: no deletion).

The domain for this dataset is "<http://purl.org/thecliff/tcl/>". One or more people from the Tcl/Tk community will be given PURL creation access eventually, but for now I ask to postpone this until the bigger picture becomes clearer and we can agree on guidelines on when to create a new PURL and on how to choose names. Extreme restraint is needed, because there is no way to get rid of a PURL once defined. Basically, every Tcl-related dataset will be given a PURL in the /tcl/ domain.

PURL ownership is simpler. If Scriptics maintains a "patches" dataset, then it is only logical to give Scriptics the tools to decide where this dataset will be located. Every PURL can have different owners, though it may be a good idea to give a few well-known people and organizations in the Tcl community access and let them adjust resources. This is likely to happen rarely. Dataset servers tend to stay in one place. Though I'm an owner now, future owners can take me off the list later.

The differences between PURLs and domain names are:

- ownership/creation of PURLs can be distributed among many
- changes to a PURL are instant, there is no caching
- every PURL can redirect to a different site

PURLs are what make it possible to say "THE dataset name of the Tcl-URL! archive is: tcl/url". Just add a standard prefix, and you'll be able to locate its associated dataset server. Append ".html", and you'll be going to a page where all the details of that dataset can be viewed.

For now, we're aiming for the following datasets:

tcl/url	Archive of weekly Tcl-URL! posts (JC)
tcl/jumbo	Alex's resource mix
tcl/patches	Patches, maintained by Scriptics

Note that the actual *location* of datasets is irrelevant. There are always master copies (in this first approach), but we don't care where.

Close your eyes. How many datasets will there be a year from now? Answer: it doesn't matter, this can scale and is location independent.

BECOMING FUTURE-PROOF

So what *is* a "dataset"? Well, that's something I have been staying away from - because I hope to deal with this in a very generic manner.

A dataset is a collection of entries, each of which can be uniquely identified with some string (a record ID, a key field, whatever). There is no inherent ordering, though keys may imply an obvious/natural one. X

Each entry also has a "modification timestamp" as attribute.

The content of entries is unspecified. Anything goes.

A "dataset server" is what you need to be able to define such datasets in the PURL context described above. It is currently defined as a very simple HTTP request/reply service. Requests use the CGI approach of adding "?..." to the dataset URL to specify the exact request.

Replies are based on the XML (extensible markup language) notation. When you talk to a dataset server, it responds in XML. The protocol of these responses is standardized so that all dataset servers act the same, though the data itself, and their structure, will vary widely. XML is a tree-structured notation which can accommodate all formats. And the way things are moving now, XML is set to sweep the net, IMO.

XML is what will make dataset servers and clients "future proof". There is a very rudimentary server, written in Tcl, running on tclhttpd, and serving the Tcl-URL! dataset right now. It's just over 100 lines. Ten years ago, it might have been written in C (you tell me how many lines), ten years from now it can be written with a tool that doesn't exist yet. XML is transport- / system- / tool- / language- / platform-independent.

Note that XML strictly defines what is exchanged between dataset servers and future dataset client(s), but that the storage of the dataset itself is unspecified. Each one can be different. Today's datasets often are.

So what does a dataset server *do* exactly? Well, not that much - just enough to make the whole system work. A dataset server must be able to tell clients what has changed since the last time they checked. It does this - in response to an HTTP request - by returning a list of commands which the client needs to perform to update itself.

Here's a hypothetical transcript:

```
Client: hi, I'm new, what do you have to offer?  
Server: ok, start with a clean slate, then add [these entries]  
Server: the last change was on 1998/10/20 12:34:56, by the way  
... time passes ...  
Client: hi again, anything happened since 1998/10/20 12:34:56?
```

```
Server: oh yes, lots: add [this] as X, delete Y, replace Z with [this]
Server: FYI, the last change was on 1998/10/27 01:23:45
... time passes ...
Client: what's up, doc? I last spoke to you on 1998/10/27 01:23:45
Server: it's been a bit quiet, nothing has happened since then
... and so on ...
```

MANAGING THE DATASETS

A lot of things depend on state changes. Servers hop from one state to the next whenever changes are made to their dataset. There are many ways in which dataset editing could be implemented. For now, I manually apply changes to the tcl/url dataset, since those are always only additions. But that will be extended soon.

One could envision Tk-based utilities which talk to a more sophisticated dataset server, identifying themselves, and applying changes from a remote location. That could be a single maintainer, or a team of people who have agreed to maintain a dataset, or eventually even an way to let anyone submit/maintain information. We'll come with many solutions.

A simpler approach, is to let the server generate HTML forms, and to let those forms interface with the server, using a CGI process or otherwise.

SERVER PROTOCOL

We now have independence of just about anything. The one aspect which must be clearly defined, is how a server responds to requests. With this standard, servers and clients can be implemented. The protocol used is new, extensible, and maximally decoupled from content. It is also minimalistic. At least, that's what I've tried to accomplish.

In other words: dataset servers can remain simple, and will handle just about any type of data we come up with later.

There will no doubt be more extensions to the protocol in the future.

Let me translate the above transcript into what goes over the wire:

```
Client: fetch http://purl.org/thecliff/tcl/url?since=0

Server: returns the following, as mime type "text/xml":
<update last="19981020123456000">
    <reset/>
    <add id=... modified=...>
        ...
    </add>
    ...
</update>

Client: fetch http://purl.org/thecliff/tcl/url?since=19981020123456000

Server: returns the following, as mime type "text/xml":
<update last="19981027012345000">
    <add id="X" modified=...>
        ...
    </add>
    <delete id="Y" modified=.../>
    <replace id="Z" modified=...>
        ...
    </replace>
</update>

Client: fetch http://purl.org/thecliff/tcl/url?since=19981027012345000
```

Server: returns the following, as mime type "text/xml":
<update last="19981027012345000"/>

Note that everything between <add> and </add> tags is the contents of that entry, expressed in well-formed XML notation. You can examine the real output to see what it can look like. It all resembles HTML, except that the tags do not specify formatting but datastructure.

Note also that the dataset server doesn't care *what* the data is. It is a rocket booster, with a payload (the data) - any payload that fits.

And that's all there is to it.

I will be adding the "modified" attribute to the XML output real soon (this change will make some neat change propagation tricks possible).

If anyone wants to experiment with this stuff on the client side, feel free to use the tcl/url dataset server. I expect Steve Ball's TclExpat and other XML parsing tools to come in quite handy to grab the contents and store it on your site. I hope to work on that a few weeks from now.

I also intend to add a larger dataset server, one which changes a few times a day, so that it will be more useful to see change propagation in action. If it turns out to be simple, maybe later this week.

Ok, shoot :)

-- JC